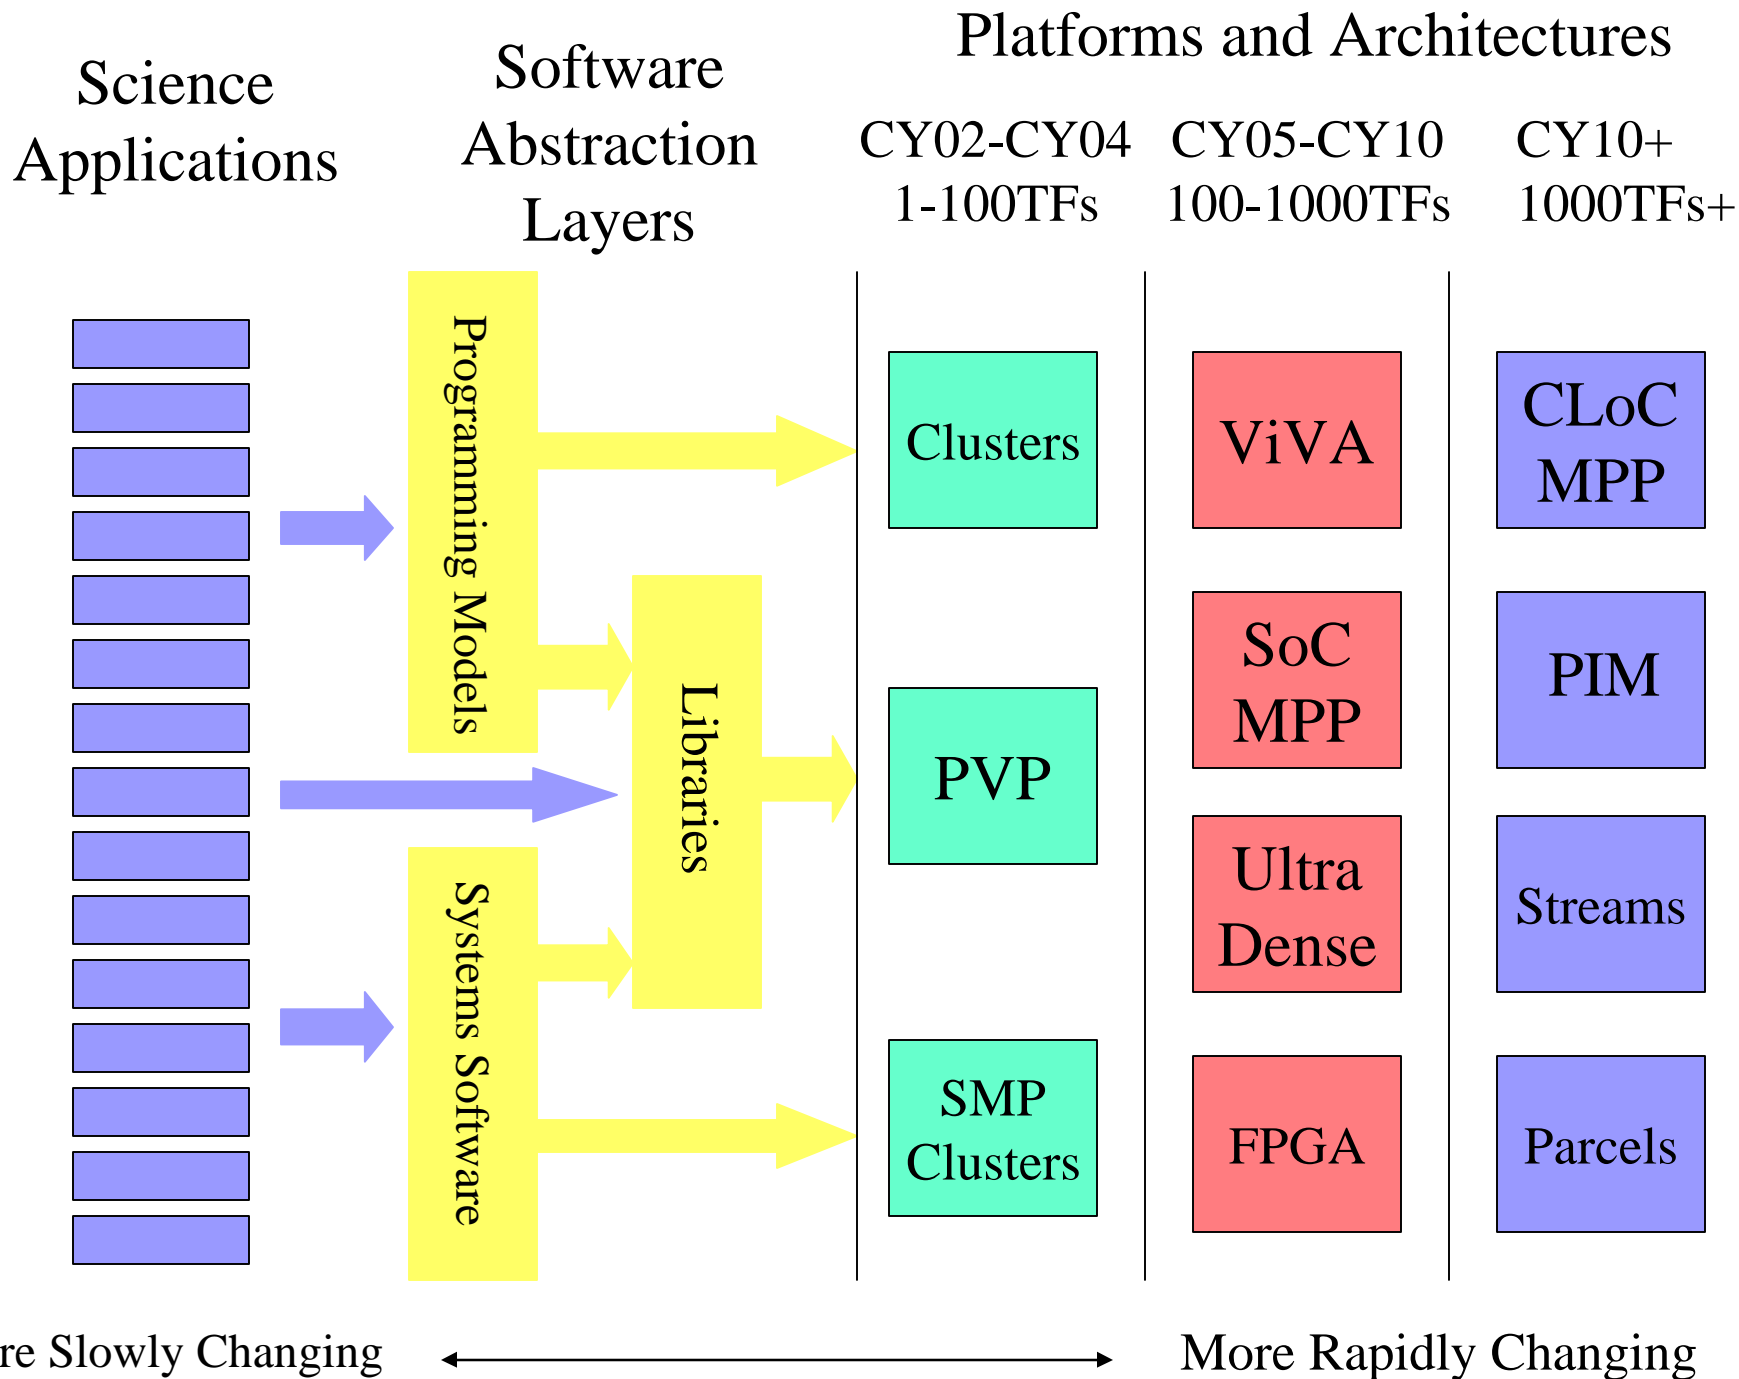


Opportunities and Challenges: Hardware and Applications

Rick Stevens

Argonne National Laboratory

University of Chicago



How Much do Apps Developers Need to Know about Architecture?

- Nothing under the assumption that:
 - PM is performance transparent
 - Architecture is capable of contributing little additional performance via the programming model
- Everything under the assumption that:
 - PM is not performance transparent
 - Architecture is not well served by the PM
- Nothing under the assumption that:
 - Real accomplishments in computational science (field) require many years to achieve
- Everything under the assumption that:
 - Large runs are rare

Mo'Betta Science

- Which will yield more/better science
 - Driving up performance?
 - Driving down cost ?
- Are we really doing science or are we doing computational stuntology ?
 - How many of the largest runs are repeated with variations in parameters ?
- Computing resources need to be so freely available that we can waste them in pursuit of crazy ideas
 - On-demand computing will dramatically improve scientific productivity
- Users need to be able to try experiments without re-coding their tools
 - Interactive applications accelerators

Petascale Problem Solving

- Most computational science is coded in Matlab, Mathematica, Maple, etc.
 - Problem is moving from high-level tools to low-level tools to run on large-scale machines
- The biggest performance GAP is the problem that is never run due to cost/time of translating from high-level ideas to scalable code
 - Due to scarcity of resource we run only the most conservative problems on the largest machines
- We could be trying building architectures that will enable interactive analysis at Petascale
 - Real-time data analysis and hypothesis formation

Architecture Futures in a Nutshell

- Sustainability
 - Need long term fountain of new architectures
- Abundance
 - Successful architectures need to become ubiquitous
- Affordability
 - The building blocks for future large-scale machines need to be the same as those for small and medium scale machines
- Scalability
 - Groups should be able to play time-space tradeoffs
- Interactivity
 - Future machines need to become more interactive